

DIAL ACCESS STACK ARCHITECTURE

5 This application is continuation of prior U.S. Serial No. 09/107,187, filed June 29, 1998.

BACKGROUND OF THE INVENTION

10 This invention relates to network access servers and more particularly to a novel dial access stack architecture.

A Network Access Server (NAS) is used for processing multiple fax, analog modem, digital data or other types of calls sent over a Public Service Telephone Network (PSTN) or any other type of communication system. The NAS includes T1, E1, T3 and/or E3 line interfaces that send and receive information over the PSTN.

15 Controllers, framers and modem modules in the NAS convert channel data from the line interface units into digital packets. The packets are sent from the modems over a backplane to router circuitry in the NAS that sends the packets out a packet based network over a LAN or WAN port.

20 As Internet traffic increases, there is a need to increase the number of communication channels that the NAS can handle at the same time. The prior solution for increasing NAS call processing capacity was to simply increase the number of line interface units, framers and modem modules in the NAS chassis. However, NAS capacity is limited to the physical number of modules that can be supported in one NAS box. Processing capacity is also limited by the bandwidth of
25 the buses used in a NAS backplane for sending data between the different NAS processing modules. Thus current NASs have limited scalability and can only process information from a limited number of communication channels.

30 The individual line interface units and other processing modules typically communicate to each other using a proprietary communication protocol. A NAS therefore cannot be easily upgraded or interchanged with modules used in other NAS architectures or incorporating different processing technology. All processing modules must also be compatible with the physical board sizes and interfaces used for connecting the modules to a NAS backplane. These restrictions also make it difficult to upgrade NASs or increase the communication links the NAS can process.

Current NAS architectures provide little or no fault tolerance against failures that occur in the field. Upon encountering a failure, field service engineers typically swap out the entire NAS box. For example, when a single modem module in the NAS fails, the entire NAS box is turned off and the modem card replaced. When the NAS is shut down, every call coming into the NAS is disrupted. Because the NAS handles a large number of calls at the same time, any failure, no matter how small, disrupts a large number of telephone calls.

Some NAS architectures break the NAS system into many very small subsystem cards. When a failure occurs, the whole subsystem card is decommissioned and manually swapped by an operator with a standby subsystem card at a later time. Even if a subsystem is partially operational, it is fully decommissioned if a failure is detected. To reduce the effects of failures, redundant cards are placed in the NAS chassis. However, the redundant cards take up space in the NAS chassis and require additional power and interconnectivity that further reduce NAS scalability.

Accordingly, a need remains for a network access server architecture that is more scalable and more easily upgradeable while at the same time being more fault tolerant.

SUMMARY OF THE INVENTION

The Dial Access Stacking Architecture (DASA) provides scalability and resiliency to fault conditions and can easily aggregate and integrate new access media. Applications such as voice, video and multicasting can be seamlessly added. The DASA architecture can scale from hundreds to thousands of ports to optimize performance. System redundancy avoids any single point of failure.

The DASA includes a stack of network access servers each independently processing information for communication links established over a public telephone network. A primary interconnect couples the stack of network access servers together through a primary network. A routing engine is coupled through the primary interconnect to the stack of network access servers. The network access servers, primary interconnect and the routing engine, in one embodiment, are all independently operating stand alone systems. The primary interconnect comprises a packet-based network switch that allows pairs of the network access servers to communicate with each other in parallel while one of the other network access servers transfers

information to the routing engine. It also provides an adequate buffer to hold packets for automatic re-transmission when the recipient is busy. As such the transmitting entity (network access server or routing engine) never has to wait. This in turn reduces the probability that the connected entities are busy. A secondary interconnect couples
5 the stack of network access servers together through a secondary network that operates independently of the primary network.

A system controller is used to monitor, configure and debug the other DASA components. The system controller independently accesses the stack of network access servers through the primary and secondary interconnect. A serial bus is also
10 coupled between the network access stack and the system controller for system debugging.

The DASA is used with a Stack Group Bidding Protocol (SGBP) to implement a large multi-link dial pool that is multiple times larger than a single NAS can support. Members of a stack group are established from the multiple network
15 access servers. Multiple communication links from one site are then established to the stack group members that operate together as a multilink bundle. The stack group members upon establishing the communication links send bid requests for mastership of the multilink bundle. One of the stack group members making a highest bid is assigned as a bundle master. The communication links in the bundle are sent to the
20 bundle master through the primary or secondary interconnect and the multilink session is conducted with the bundle master through the interconnect independently of the router. This reduces bottlenecks that could occur in architectures where all communications are required to go through the router or a central processing unit.

Unlike proprietary scaling solutions that use customized hardware, DASA has
25 no theoretical constraint on scaling. The only limiting factor is the processing power of each component. The DASA allows integration with a wide selection of readily available commercial components. Thus, as technology advances and more powerful or lower cost components become available, the new components can be easily integrated into the DASA system.

30 The foregoing and other objects, features and advantages of the invention will become more readily apparent from the following detailed description of a preferred embodiment of the invention, which proceeds with reference to the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a Dial Access Stack Architecture (DASA) according to the invention.

FIG. 2 is a detailed block diagram of a Network Access Server (NAS) stack
5 used in the DASA shown in FIG. 1.

FIG. 3 is a detailed block diagram of routing engines used in the DASA shown in FIG. 1.

FIG. 4 is a detailed logic diagram showing how a primary or secondary interconnect are used in the DASA shown in FIG. 1.

10 FIG. 5 is a block diagram showing how the DASA is used in a multilink point to point protocol session.

DETAILED DESCRIPTION

Referring to FIG. 1, a DASA 12 comprises five major components or sub-
15 systems. A Network Access Server (NAS) stack 16 is connected to a Public Service Telephone Network (PSTN) 14. A primary interconnect 18 and a secondary interconnect 20 are each coupled between the NAS stack 16 and routing engines 22. The routing engines 22 are coupled to Internet 26. A system controller 24 is connected directly to the NAS stack 16 through a serial interconnect 23 and indirectly
20 to the NAS stack 16 through interconnects 21 connected to primary interconnect 18 and secondary interconnect 20. An interconnect 19 is connected from the system controller 24 to routing engines 22.

It should be understood that the DASA 12 can be used in any multiservice application. For example, the stack 16 can receive communication links via cable
25 modems, XDSL, PSTN, frame relay or any other type of communication system. For clarity, the invention is described in relation to PSTN but is not limited in scope only to PSTN networks. The scope of the invention covers a DASA that concentrates multiple access points for any system.

The NAS stack 16 aggregates analog and digital calls on dial links 28
30 established over the PSTN 14 from dial up clients 30. Packet streams extracted from the dial up client calls are switched between individual network access servers and the routing engines 22 by the primary interconnect 18. In the event of a failure in the primary interconnect 18, traffic is switched between DASA components using the

secondary interconnect 20. The converse is true for packet traffic from Internet 26 to the dial up clients 30. The packets coming from Internet 26 are switched from the routing engines 22 to the NAS stack 16 across the primary interconnect 18. In the event of a failure in the primary interconnect 18, secondary interconnect 20 is used.

5 The system controller 24 manages and monitors the components in DASA stack 12 and also serves as the primary Network Timer Protocol (NTP) clock source for the members of the DASA stack 12. A modem 27 is used to send management data extracted from the DASA components to a central network manager (not shown). The system controller 24 is used for out of band access to all DASA components. The
10 system controller features triple redundant path to each component: primary interconnect, secondary interconnect, serial connection.

 All the components used in the DASA 12 are commercially available from Cisco Systems, Inc. 170 West Tasman Drive, San Jose, CA 95134-1706 or from other manufactures. In one implementation, the network access servers in NAS stack 16 are
15 Cisco Model No. AS5300 network access systems that each support multiple PRI (T1/E1) lines 28, a 100BaseT full duplex Ethernet 34 and a 10BaseT Ethernet 36. The primary interconnect 18 is a Cisco Catalyst Model No. 5002 100BaseT Ethernet switch and the secondary interconnect 20 is a Cisco Model No. 7206 10BaseT Ethernet switch. The routing engines 22 are Cisco Model No. 7206 or 3640 routers.
20 The system controller 24 is a Cisco Model No. AS3640 router. The system controller 24 and the routing engines 22 can also be collapsed into one Cisco AS3640 router with extra ports for console management.

 It should be understood, however, that the DASA 12 could incorporate different components from different network system manufacturers. For example, the
25 NASs in NAS stack 16, the switches in the primary and secondary interconnects 18 and 20, the routers in routing engines 22 and the system controller 24 are all interchangeable with any other system that can perform similar functions. This is because industrial standardized interfaces and protocols are used between all of the different DASA components.

30 The DASA 12 has no commercial limits on stacking size. The routing engines 22 can incorporate new router technology that operates at least at the speed of the internet 26. Thus, the routing engines will never create a bottleneck for transferring information over the internet.

The interconnects 18 and 20 can also be upgraded with new network switches with more ports. Alternatively, two or more switches can be connected together to increase the total number of ports available for connecting NAS's together. The hierarchy of interconnects can vary for different DASA arrangements. For example, one interconnect can be used to connect two other interconnects to the routing engines 22. The first interconnect that connects the other two interconnects together should be able to operate at a higher data rate to prevent bottlenecks between the other two interconnects and the router.

Referring to FIG. 2, the NAS stack 16 comprises banks of NASs 32 organized into a coherent sub-system. Each NAS 32 has multiple Wide Area Network (WAN) dial links 28 and ports for two LANs 34 and 36. The LAN 34 is tied from the NASs 32 to the primary interconnect 18, while LAN 36 is tied from the NASs 32 to the secondary interconnect 20.

The NASs 32 each include a console port that connects through a serial interconnect 23 to system controller 24 (FIG. 1). In one embodiment, the console port is a serial RS232 port used by the system controller 24 to configure the NAS stack 16 or as an emergency link for remote diagnostic and trouble shooting. The NASs 32 may also have additional ports to provide more redundancy, load sharing, or special signaling segregation. The NAS ingress ports connected to dial links 28 are typically implemented as a telephone company trunk such as PRI, BRI, or xDSL. Primary Rate Interface (PRI), Basic Rate Interface (BRI), Digital Subscriber Loop (XDSL) There are many types of DSL, xDSL means the general class of DSL. Another implementation uses a cable modem. The NASs 32 can also be implemented using multiple physical entities. For example, the NASs 32 can comprise a standard router joined to a xDSL concentrator via an ATM link.

As NASs 32 are added in the stack 16 to increase dial port density, the relative processing power available for each dial link 28 stays constant. For multi-link applications, such as those using a Stack Group Bidding Protocol (SGBP) as described below in FIG. 5, processing incurs minor overhead. If other components in DASA 12, such as the primary interconnect 18, secondary interconnect 20 and routing engines 22 are scalable enough, there are no practical limits to the stack size and the number of dial links 28 that can be supported by NAS stack 16.

The primary and secondary interconnects 18 and 20 typically comprise network switches that interconnect the NASs 32 and routing engines 22 together. The primary interconnect 18 is the primary data path between the components in DASA system 12. The secondary interconnect 20 provides redundancy in the event of a failure in primary interconnect 18. The backup path established by secondary interconnect 20 is configured for automatic and fast convergence with negligible service disruption if the primary path established by primary interconnect 18 fails. The secondary interconnect is optional and can be omitted if redundancy is not an important factor in the system.

Referring to FIG. 3, any router can be used for the routing engines 22. For redundancy and load balancing, 2 routing engines 38 and 40 are used to back up each other in case of a failure. The routing engines 38 and 40 are each coupled to both the primary interconnect 18 and the secondary interconnect 20. If the primary routing engine 38 fails, routing engine 40 is automatically activated to support either the primary interconnect 18 or the secondary interconnect 20 for transferring packets between the NAS stack 16 and Internet 26. The routing engines 38 and 40 include multiple egress ports for establishing networks 34 and 36. A port establishes connection 39 to Internet 26. The egress ports are typically implemented in high speed LAN or WAN interfaces such as 100BaseT, ATM or Optical Fiber.

FIG. 4 is a detailed logic diagram showing how the primary and secondary interconnects 18 and 20 operate. The primary and secondary interconnects 18 and 20 each consist of many high speed ports compatible with the interfaces of other components in DASA 12. For optimal performance, the interconnects 18 and 20 support simultaneous non-blocking data transfer of all the interconnections through a network switch that includes switching circuits 50, 52 and 54. There are many switching circuits in both the primary and secondary interconnects 18 and 20. For clarity, only three switching circuits are shown in FIG. 4. The switching circuits 50, 52 and 54 are each interconnected to each one of the components in the DASA 12 at opposite ends 58 and 60. A controller 56 configures any one of the switches 50, 52 and 54 to connect one of the DASA components at one end 58 to another DASA component at the opposite end 60.

A common buffer 55 is connected to controller 56. The common buffer 55 temporarily stores packets sent to NAS's that are currently busy. Thus, one NAS 32

can continue to send packets for later processing by the busy NAS 32. The common buffer 55 therefore reduces bottlenecks and delays that could occur when sending data between NAS's 32.

The interconnects 18 and 20 allow different NASs 32 and routing engines 38 and 40 to connect to each other and transfer information in parallel. For example, a first and second NAS can transfer information between each other at the same time that a third and fourth NAS 32 are transferring information with each other and a fifth NAS 32 is transferring information with one of the routers 38 or 40. This is a significant improvement over previous NAS implementations that require all data traffic to go through the router or a central processing entity. A router only has a single processor for transferring packet traffic and must route data traffic serially as compared to the parallel packet traffic provided by switches 50, 52 and 54.

The network switches 18 and 20 are designed to connect together a large number of devices at the same time. Therefore, the DASA is easily scaled to add any number of NASs 32 or routing engines 22 to the DASA 12. The interconnect switches 18 and 20 in one embodiment use LAN interfaces that are the common interface used by the different NASs 32 and routing engines 22. The LAN interfaces allow almost any commercially available NAS or router can be incorporated into DASA 12. The DASA 12 also allows the primary and secondary interconnects 18 and 20 to be easily upgraded with new switching technology that may provide more ports with faster LAN connections.

The primary and secondary interconnects 18 and 20 operate essentially as a multiplexer for the routing engines 22. For example, without the primary interconnect 18, the routing engines 22 would need one port for each NAS 32. The number of ports on the routing engines 22 would then limit the scalability of the DASA 12. However, the primary and secondary interconnects 18 and 20 effectively multiplex all the NASs 32 through a single router port. Thus, the routing engines 22 can support any number of NASs 32. The interconnects 18 and 20 also allow the NASs and routing engines 22 to be physically separated from each other. Thus, all the components in DASA 12 do not have to be located in the same chassis. This eliminates mechanical limitations, such as cooling requirements and box size, that could limit the scalability of the DASA 12.

Any protocol such as an Open Shortest Path First (OSPF) protocol can be used as the routing protocol that binds the whole DASA system 12 together. For example, an Enhanced Interior Gateway Routing Protocol (EIGRP), developed by Cisco Systems, Inc., selects the best path for packet forwarding based on bandwidth, delay and loading. Routing selection within the DASA 12 can be controlled by manipulating the delay and bandwidth of the different LAN connections. To select one path over the other, the bandwidth is set higher for that LAN interconnection to reduce delay. EIGRP routes data traffic over the higher bandwidth primary path under normal operation and over the secondary path in the event of failure in the primary path.

Multilink Sessions

Referring to FIG. 5, for multi-link calls, the NAS stack 16 is configured into a stack group using a Stack Group Bidding Protocol (SGBP). SGBP is a protocol that binds selectable NAS 32 in the NAS stack 22 into a single logical access server. Multi-link calls are then distributed across different chassis within the NAS stack 16. This arrangement encompasses all the multi-link functions, packet fragmentation and packet reassembly within the NAS stack 16. The SGBP is described in detail in co-pending patent application Ser. No. 08/846,788 entitled: DYNAMIC BIDDING PROTOCOL FOR CONDUCTING MULTILINK SESSIONS THROUGHOUT DIFFERENT PHYSICAL TERMINATION POINTS filed April 30, 1997 which is herein incorporated by reference.

In a multi-link call, two streams of packet fragments can travel over two dial links 44 and 46. These links may be terminated by two different NASs 32A and 32B. The packet fragments from the dial links 44 and 46 must be reassembled before being forwarded to the routing engines 22. The NAS 32A that terminated the first call is configured by the SGBP to reassemble data streams from both dial links 44 and 46, while the NAS 32B that terminated the second link 44 forwards its data stream to the first NAS 32A. The reassembled stream 48 is then forwarded across one of the primary or secondary interconnect 18 or 20, respectively, avoiding packet fragment processing at the routing engines 22 or at the system controller 24.

In the SGBP, a pipe 42 is established between stack group members 32A and 32B. When a task or event occurs, such as establishment of communication links 44

and 46, the stack group members 32A and 32B conduct a bidding process. During the bidding process, each stack group member 32A and 32B bids for the event. The bidding process uses weighting criteria that varies dynamically depending on the computational status of the NASs 32 at the time the bidding is initiated. The value bid
5 by each NAS 32 is weighted according to whether the NAS 32 is already controlling or processing similar events, network locality of the NAS 32 in relation to the event, CPU capacity of the NAS 32, current loading of the NAS 32, manual override values and an offload criteria that indicates the NAS 32 making the bid does not want to process the event. The event is allocated to the NAS 32 making the highest bid.

10 In one specific application, a multichassis multilink PPP (MLP) protocol utilizes SGBP to conduct multilink PPP sessions for links either originating or terminating on different physical systems. The SGBP establishes the different NASs 32, say NAS 32A and 32B, as stack group members that serve as termination points for an MLP bundle. The SGBP bidding scheme then establishes one of the NAS
15 members, say NAS 32A as master of the MLP bundle. Upon receiving an incoming call upon which MLP has been negotiated, the stack group member 32B determines if a bundle already exists within the stack group where the link can be added. Since a bundle exists under the mastership of another stack group member 32A, the link is forwarded to that master as the "tunneled" link 42. The MLP protocol is well known
20 to those skilled in the art, and is therefore, not described in further detail.

A L2F forwarding protocol is used in combination with multichassis MLP to forward the link 42 from NAS 32B to NAS 32A and offers location transparency. The L2F forwarding protocol is encapsulated around PPP sessions and then sent over a tunnel to the physical NAS 32A conducting the multilink session. The L2F
25 forwarding protocol is described in co-pending patent application Ser. No. 08/687,973 entitled: "Virtual Dial-Up Protocol for Network Communications" filed July 7, 1996 which is herein incorporated by reference.

Without primary or secondary interconnect 18 or 20, all NASs 32 would have to contend the same network for communicating with each other and also with the
30 routing engines 22. The primary and secondary interconnects 18 and 20 allow multiple pairs of NASs 32 to communicate to each other at the same time. Even more significant, the primary and secondary interconnects 18 and 20 allow anyone of the NASs 32 to communicate with the routing engines 22 while the other NASs 32

communicate with each other in parallel. This provides the substantial advantage of allowing the NASs 32 to perform the multilink PPP sessions independently of the data transfer process between the NAS stack 16 and routing engines 26. The transfer of information between the NAS stack 16 and the routing engines 22 is therefore more efficient, because the multilink packets have already been bundled together into one packet stream before being sent to routing engines 22.

System Controller

Referring back to FIG. 1, the system controller 24 manages and monitors the state of the dial network 14 and the individual components of the DASA stack 12. The controller 24 also serves as the primary NTP clock source for other members of the DASA stack 12. NTP is used to keep a consistent overall view of the entire DASA system 12 for keeping the time stamped log, error messages, accounting, authentication and authorization records synchronized. The system controller 24 is used as a terminal server for console access to all components in the DASA stack 12 and performs remote out-of-band management and diagnostics. The system controller 24 monitors the traffic through the DASA 12 while bypassing the routing engines 22. Without the system controller 24, the routing engines 22 would have to be used for communicating all management information to each one of the components in DASA 12. The system controller 24 through interconnects 21 access management and diagnostic data from the DASA components and combines the data together in local memory. The stored data is then sent via a bulk data transfer through modem 27 to the network management station.

For regulatory or technical reason, some telephone companies require separate network paths for network management traffic and for data payload traffic. The system controller 24 provides these separate network paths 21 for management traffic. The serial connection 23 is coupled to each NAS 32 in the NAS stack 16 for system debugging. The NAS stack 16 would typically not be accessible if the primary and secondary interconnects 18 and 20 both went down. The serial connection 23 provides a separate out of band link to NAS stack 16 for debugging the DASA system 12.

Having described and illustrated the principles of the invention in a preferred embodiment thereof, it should be apparent that the invention can be modified in

arrangement and detail without departing from such principles. I claim all modifications and variation coming within the spirit and scope of the following claims.